

CLAIMS

What is claimed is:

1. In a network device operative to control data flows transmitted between hosts connected to a computer network, wherein at least some of the hosts employ slow-start mechanisms, a method comprising
 - estimating the initial rate demand for a data flow between a first host and a second host;
 - allocating a fraction of the initial rate demand for the flow;
 - maintaining a count of the packets associated with the flow; and
 - increasing the fraction of the initial rate demand allocated to the flow as the count crosses at least one threshold.
2. The method of claim 1 further comprising
 - estimating the number of bytes that the first host will transmit before achieving the initial rate demand; and
 - setting the at least one threshold based on the number of bytes in the second estimating step.
3. The method of claim 2 wherein the second estimating step comprises
 - estimating the round trip time between the first and second host; and
 - multiplying the initial demand rate associated with the data flow by the round trip time.
4. The method of claim 3 wherein the round trip time is based on an analysis of the arrival times of the handshake packets corresponding to the data flow.
5. The method of claim 1 wherein the initial rate demand is based on an analysis of the arrival times of at least one of the handshake packets corresponding to the data flow.

6. The method of claim 5 wherein the initial rate demand is based on an analysis of at least one data packet corresponding to the data flow.
7. The method of claim 1 wherein the initial rate demand is based on an analysis of at least one data packet corresponding to the data flow.
8. The method of claim 1 further comprising
 - estimating the number of packets that the first host will transmit before achieving the initial rate demand; and
 - setting the at least one threshold based on the number of packets in the second estimating step.
9. The method of claim 8 wherein the second estimating step comprises
 - estimating the round trip time between the first and second host;
 - multiplying the initial demand rate associated with the data flow by the round trip time; and
 - dividing the product of the multiplying step by an average packet size.
10. The method of claim 9 wherein the average packet size is a static parameter.
11. The method of claim 10 wherein the average packet size is a configurable parameter.
12. The method of claim 9 wherein the average packet size is a dynamic parameter that changes based on observations of the packets traversing the network device.
13. The method of claim 9 wherein the round trip time is based on an analysis of the arrival times of the handshake packets corresponding to the data flow.

14. The method of claim 8 wherein the initial rate demand is based on an analysis of the arrival times of at least one of the handshake packets corresponding to the data flow.

15. The method of claim 14 wherein the initial rate demand is based on an analysis of at least one data packet corresponding to the data flow.

16. The method of claim 8 wherein the initial rate demand is based on an analysis of at least one data packet corresponding to the data flow.

17. The method of claim 1 further comprising
 monitoring for at least one indication that the sending host has re-initiated the slow start mechanism for the data flow;
 upon detection of at least one of the indications,
 resetting the count of the packets for the flow; and
 repeating the allocating, maintaining and increasing steps.

18. The method of claim 17 wherein the monitoring step comprise
 determining whether at least one data packet corresponding to the data flow is a re-transmission of a previous packet.

19. The method of claim 18 wherein the monitoring step further comprises
 determining whether the re-transmitted packet arrived a threshold period of time after the last packet corresponding to the data flow.

20. The method of claim 17 wherein the monitoring step comprise
 determining whether the packet arrived a threshold period of time after the last packet corresponding to the data flow.

21. An apparatus facilitating control data flows transmitted between hosts connected to a computer network, wherein at least some of the hosts employ slow-start mechanisms, comprising

- a packet processor operative to

- detect a data flow in network traffic traversing a communications path;

- maintain a count of the packets associated with the data flow;

- a path rate detection module operative to

- estimate the initial rate demand for a data flow;

- estimate, for the data flow, the number of bytes that a sending host will transmit to a receiving host before achieving the initial rate demand;

- a bandwidth allocation module operative to

- allocate bandwidth to the data flow based in part on a target rate associated with the data flow; and

wherein the apparatus is operative to

- set the initial target rate for the data flow as a fraction of the initial rate demand for the flow; and

- increase the target rate associated with the data flow as the count of bytes crosses a threshold value.

22. The apparatus of claim 21 further comprising

- a traffic classification database including at least one traffic class, at least one attribute defining the at least one traffic class, and at least one bandwidth utilization control corresponding to the at least one traffic class, wherein the traffic classification is operative to

- compare attributes stored in association with traffic class identifiers to attributes of the data flow to identify a traffic class that corresponds to the data flow;

- associate at least one bandwidth utilization control to the data flow based on the identified traffic class; and

wherein the bandwidth allocation module is operative to allocate bandwidth based, at least in part, on the target rate, and the at least one bandwidth utilization control, associated with the data flow.

23. The apparatus of claim 21 wherein the packet processor is further operative to parse at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type a pointer to an application-specific attribute.

24. The apparatus of claim 21 wherein the packet processor is further operative to parse at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type a pointer to an application-specific attribute.

25. The apparatus of claim 24 wherein the traffic classification database is operative to match the flow specification to a plurality of traffic classes, each of the traffic classes defined by one or more matching attributes; and thereupon, having found a matching traffic class in the matching step, associate said flow specification with a traffic class from the plurality of traffic classes.

26. The apparatus of claim 21 wherein the apparatus is further operative to monitor for at least one indication that the sending host has re-initiated the slow start mechanism for the data flow; and upon detection of at least one of the indications, reset the count of bytes for the flow; and

reset the target rate for the data flow to the initial target rate.

27. The apparatus of claim 21 wherein, to estimate the number of bytes that a sending host will transmit before achieving the initial rate demand, the path rate detection module is operative to

- estimate the round trip time between the sending and receiving host; and
- multiply the initial demand rate associated with the data flow by the round trip time.

28. The apparatus of claim 27 wherein the round trip time is based on an analysis of the arrival times of the handshake packets corresponding to the data flow.

29. The apparatus of claim 27 wherein the initial rate demand is based on an analysis of the arrival times of at least one of the handshake packets corresponding to the data flow.

30. The apparatus of claim 29 wherein the initial rate demand is based on an analysis of at least one data packet corresponding to the data flow.

31. The apparatus of claim 21 wherein the initial rate demand is based on an analysis of at least one data packet corresponding to the data flow.

32. An apparatus facilitating control data flows transmitted between hosts connected to a computer network, wherein at least some of the hosts employ slow-start mechanisms, comprising

- a packet processor operative to

- detect a data flow in network traffic traversing a communications path;
 - maintain a count of the packets associated with the data flow;

- a path rate detection module operative to

- estimate the initial rate demand for a data flow;

estimate, for the data flow, the number of packets that a sending host will transmit before achieving the initial rate demand;

a bandwidth allocation module operative to

allocate bandwidth to the data flow based in part on a target rate associated with the data flow; and

wherein the apparatus is operative to

set the initial target rate for the data flow as a fraction of the initial rate demand for the flow; and

increase the target rate associated with the data flow as the count of packets crosses a threshold value.

33. The apparatus of claim 32 wherein the apparatus is further operative to

monitor for at least one indication that the sending host has re-initiated the slow start mechanism for the data flow; and

upon detection of at least one of the indications,

reset the count of packets for the flow; and

reset the target rate for the data flow to the initial target rate.

34. The apparatus of claim 32 wherein, to estimate the number of bytes that a sending host will transmit before achieving the initial rate demand, the path rate detection module is operative to

estimate the round trip time between the sending and receiving host;

multiply the initial demand rate associated with the data flow by the round trip time; and

divide the product of the multiplying step by an average packet size.

35. The apparatus of claim 34 wherein the average packet size is a static parameter.

36. The apparatus of claim b3 wherein the average packet size is a configurable parameter.

37. The apparatus of claim 34 wherein the average packet size is a dynamic parameter that changes based on observations of the packets traversing the network device.

38. The apparatus of claim 34 wherein the round trip time is based on an analysis of the arrival times of the handshake packets corresponding to the data flow.

39. The apparatus of claim 32 wherein the initial rate demand is based on an analysis of the arrival times of at least one of the handshake packets corresponding to the data flow.

40. The apparatus of claim 39 wherein the initial rate demand is based on an analysis of at least one data packet corresponding to the data flow.